Haswell Processors

Microarchitecture

The Intel Haswell processor incorporated into the Pleiades supercomputer is the 12-core E5-2680v3 model. Haswell uses the 22 nanometer (nm) transistors that were introduced with Ivy Bridge processor, but has a newer microarchitecture that focuses on lower power consumption and higher efficiency.

Instruction Sets

Like the Sandy Bridge and Ivy Bridge processors, Haswell supports single instruction, multiple data (SIMD) instruction sets, including several generations of Streaming SIMD Extensions (SSE, SSE2, SSE3, Supplemental SSE3, and SSE4), Advanced Encryption Standard (AES), and Advanced Vector Extensions (AVX).

In addition, the AVX2 instruction set was introduced with Haswell processors.

Main Features of AVX2

Features include:

- Floating-point fused multiply-add (FMA) support, which can double the number of peak floating-point operations compared with those run without FMA. With 256-bit floating-point vector registers and two floating-point functional units, each capable of FMA, a Haswell core can deliver 16 floating-point operationsâ double the number of operations a Sandy Bridge or Ivy Bridge core can deliver.
- Integer-vector instructions support has been extended from 128-bit to 256-bit capability.

- Gather vectorization support, enhanced to enable vector elements to be loaded from non-contiguous memory locations.

Compiler Support for AVX2

AVX2 is formally supported by Intel compilers starting with version 12.1. To take advantage of AVX2, use the following compiler and options:

Compiler

You can use **`comp-intel/2012.0.032`** or later modules. We recommend using **`comp-intel/2015.0.090`**.

Options

Use either **`-xCORE-AVX2`** or **`-axCORE-AVX2`**.

Both the **`-xCORE-AVX2`** and **`-axCORE-AVX2`** options turn on **`-fma`**, which computes **`a x b + c`** in one rounding. This provides higher accuracy than **`-no-fma`**, which computes **`a x b + c`** in two steps (first **`a x b`**, then **`+ c`**) and two roundings. Turn off FMA if you want your results to match legacy ones.

TIP: An application that is compiled with AVX2 instructions can run only on Haswell and Broadwell nodes. If you want a single executable that will run on any of the Pleiades processor types, with suitable optimization to be determined at run time, you can compile your application using the option **`-O3 -ipo -axCORE-AVX2 -xSSE4.2`**.
For Fortran codes, consider using the option **`-align array32byte`** to align your vector arrays on 32-byte boundaries for best performance and reliability. This option is available in Intel compiler version 13 (for example, comp-intel/2013.1.117 and later).

AVX2 is also supported in the GNU Compiler Collection (GCC) starting with version 4.7. Use compiler options such as **`-march=core-avx-2`** or **`-mavx2`**, **`-mfma`** if you want to use AVX2.

Hyperthreading

Hyperthreading is turned ON.

Turbo Boost

Turbo Boost is turned ON.

Memory Subsystems

The memory hierarchy of Haswell is as follows:

- L1 instruction cache: 32 KB, private to each core
- L1 data cache: 32 KB, private to each core
- L2 cache: 256 KB, private to each core
- L3 cache: 30 MB, shared by 12 cores in each socket
- Memory: 64 GB per socket, total of 128 GB per node

The Haswell nodes are equipped with higher speed (2,133 MHz) DDR4 memory to provide higher memory bandwidth. There are four 2,133 MHz memory channels per socket. Each channel can be connected with up to two memory DIMMs. Of the eight memory DIMM slots for each socket, four are populated with 16-GB error correcting code (ECC)-registered DDR4 memory, for a total of 64 GB per socket. With two sockets in a node, the total memory per node is 128 GB.

Connecting the two sockets are two Intel QPI links running at a speed of 9.6 gigatransfers per second (GT/s) . Each link contains separate lanes for the two directions. The total bandwidth (2 links x 2 directions) is 38.4 GB/sec.

Network Subsystem

The Haswell and Broadwell nodes are connected to the two fabrics (ib0 and ib1) of the Pleiades InfiniBand (IB) network in a different way from the earlier processor nodes, as shown in the diagram below. Note the following differences:

- In Haswell and Broadwell node connections, each of the two PCI Express Interfacesâ one from each socketâ is connected to a separate single-port, four-lane, Fourteen Data Rate (4X FDR) host channel adapter (HCA), in a dual single-port FDR IB mezzanine card.
- In Sandy Bridge and Ivy Bridge node connections, only one of the two PCI Express Interfaces is connected to a dual-port FDR IB HCA, in a dual-port FDR IB mezzanine card.

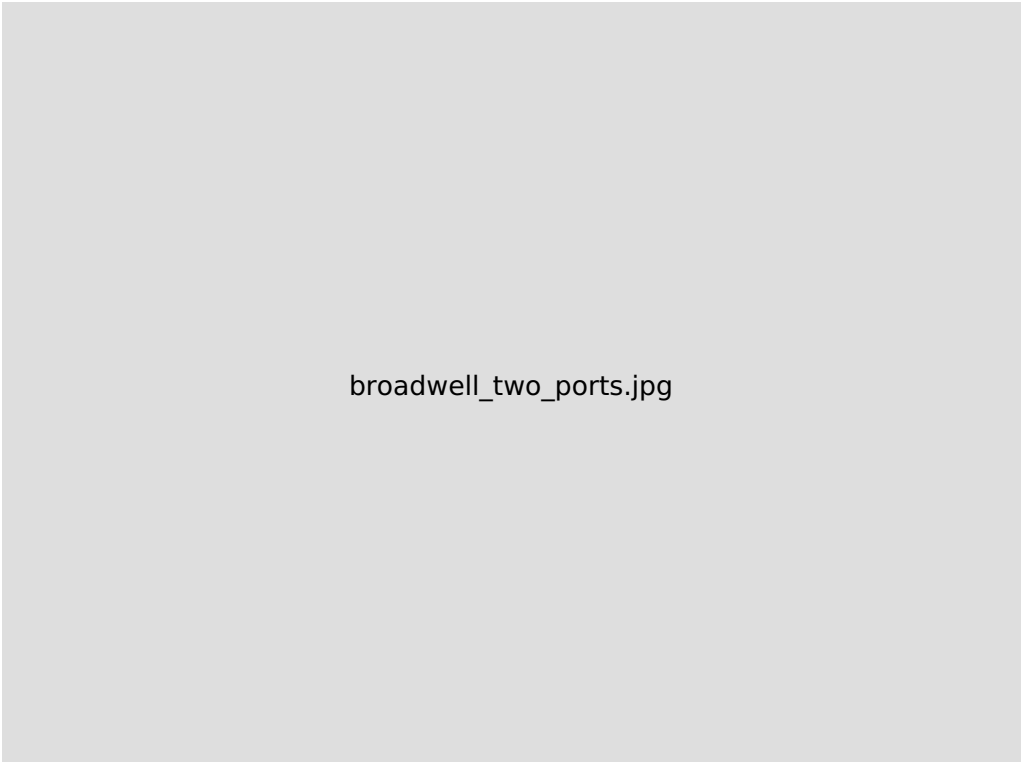

The IB mezzanine card sits on a "sister board" next to the motherboard on each node. The motherboard contains the two processor sockets. There are 18 nodes per individual rack unit. To join the ib0 fabric, the nodes are connected to two Mellanox FDR IB switches, in an ICE X IB Premium Blade. To join the ib1 fabric, another set of connections between the 18 nodes and a second Premium Blade is established.

Article ID: 492
Last updated: 23 Dec, 2020
Revision: 34
Systems Reference -> Pleiades -> Haswell Processors
https://www.nas.nasa.gov/hecc/support/kb/entry/492/